

The Application of a Course-of-Action Ontology to Support OPFOR COA Selection and Assessment

Timothy Darr, Richard Mayer and Perakath Benjamin

Knowledge-Based Systems, Inc.
1408 University Drive East
College Station, TX 77840
 {tdarr, rmayer, pbenjamin}@kbsi.com

Abstract. This paper describes the application of a course-of-action (COA) ontology to a demonstration scenario that the authors' company participated in that included a task to forecast the COAs to be executed by a simulated insurgent opposing force (OPFOR). The COA ontology includes standard decision-theoretic concepts to describe preference models from the perspective of an insurgent group for the purpose of predicting possible OPFOR COAs. The OPFOR preference structure is represented as a preference graph that visually displays the ranking of the COAs, from the perspective of the OPFOR decision maker, highlighting the most and least preferred COAs.

Keywords: Course of action planning, decision theory, utility theory, ontology, preference modeling.

1 Introduction

This paper describes the application of a course-of-action (COA) ontology [1,2] to a demonstration scenario that the authors' company participated in that included a task to forecast the COAs to be executed by a simulated insurgent opposing force (OPFOR). The COA ontology applies to the COA planning processes defined for the United States Army and Marine Corps for multiple domains, to include stability operations planning [3], counterinsurgency planning [4] and information operations planning [5]. The core ontology includes definitions of the common concepts and properties for defining COA plans, including: COAs, COA activities, COA phases, measures-of-performance (MOP) and measures-of-effectiveness (MOE).

To illustrate the use of the COA ontology, we use a scenario that is inspired by the Empire Challenge '10 (EC 10) demonstration held at Fort Huachuca in August 2010 [6-8]. In this scenario, a coalition force (CF) unit is engaged in stability operations in an area that is under contention by an insurgent OPFOR whose goals include the following:

- Short-Term: inflict damage on the CF
- Medium-Term: discourage confidence in the host nation government

- Long-Term: establish religious and cultural control

In the area of operations (AOR) under contention, there are two tribal groups whose interests conflict with each other. Tribe A is generally supportive of the CF and tribe B is generally supportive of the insurgency. These loyalties are motivated in part by a long-standing set of grievances between the tribes: tribe A and tribe B do not like each other very much, but have negotiated an uneasy truce at the moment. There are no hostilities at this time, but there is a risk that hostilities could re-emerge at any time.

Fig. 1 shows the larger operational context in which the COA ontology is used. The Hidden Enemy Network Influence Operations Map (HENIOMAP) is an application under development that is used by decision makers to assess and forecast possible COAs to aid in their own planning. The output COA ranking is an ordering of the possible COAs given as input that clearly shows the most and least preferred COAs that are consistent with a preference model. The COAs that are provided as input can be either own force or OPFOR COAs. The preference model represents the trade-offs over multiple, conflicting attributes that a decision maker employs to select the most-preferred COA to achieve their goals. The HENIOMAP ontology, of which the COA ontology described in this paper is a component, defines the domain under consideration. The HENIOMAP algorithms are used to generate the COA rankings.

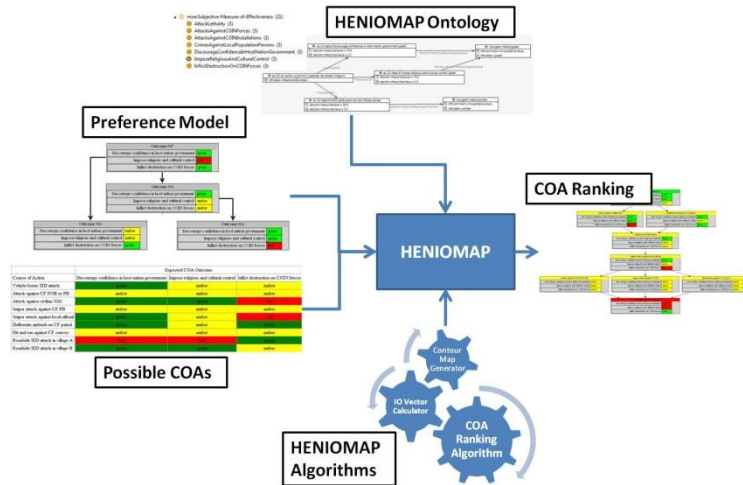


Fig. 1. Operational context for the COA ontology

Of the elements shown in **Fig. 1**, this paper describes part of the HENIOMAP ontology (COA ontology), the possible COAs, the preference model and a visualization of the COA ranking as used in the EC 10 demonstration. The details of the HENIOMAP algorithms are out of scope for this paper. Section 2 provides an overview of utility theory and a representation of some of the concepts in utility theory that are used to model COA selection and assessment problems. Section 3

provides an overview of the COA ontology that was used in the EC 10 demonstration. Section 4 presents some conclusions.

2 Utility Theory Overview

Before describing the COA ontology and how it supports the example scenario, we provide an overview of utility theory. The COA ontology models OPFOR COA selection and assessment as a multi-objective decision problem. Given a broad overall objective to be realized by a specific operation, the decision-maker must select the “best” COA to perform to achieve some objective or to identify the “best” outcome to try to achieve via a sequence of actions, where “best” is defined as an outcome that satisfactorily trades the conflicting objectives against one another from the perspective of a given decision maker.

Utility theory was originally developed in economics to measure the desirability of a good or alternative from the perspective of an agent [9]. In this model, a COA outcome replaces a "good or alternative" in the economics application, and a COA planner or decision maker replaces an "agent" in the economics application.

A utility function is given by

- $u: O \rightarrow \mathbb{R}$, where O is a COA outcome and \mathbb{R} is a real-valued number

A common form of utility function is a weighted sum of attribute values

- $u(o_i) = \sum_k w_k * a_k[o_i]$, where w_k is an attribute weight, and $a_k[o_i]$ is an attribute-value score that assigns a real-valued number to the attribute value a_k for outcome o_i . The weights of each attribute are assigned by an SME / analyst or by using an algorithm to estimate the weights
 - In this domain, the $a_k[o_i]$ is a numeric value assigned to a goal that describes a COA outcome

A preference is a relation between two outcomes such that $u(o_i) \geq u(o_j)$.

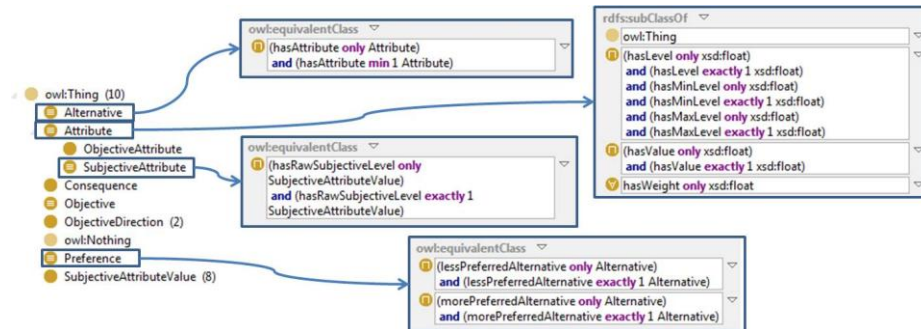


Fig. 2. Decision theory concepts used in the COA ontology

Fig. 2 shows concepts in the COA ontology that represent the elements of the utility-theoretic definition. The alternative class corresponds to the utility theory outcome and is described by a collection of attributes. The attribute class corresponds

to the utility theory attributes and are described by attribute levels, utility-theoretic values and weights. The preference class corresponds to the utility theory preference and is a relation between alternatives in which one alternative is preferred to another attribute, from the perspective of a given decision maker. In addition, a subjective attribute class is added as a subclass of attribute to represent attributes whose values are non-numeric. Extensions of these concepts for the COA ontology, along with examples, will be given in the next section in the context of the EC 10 demonstration.

3 Course-of-Action Ontology to Support EC 10

This section describes extensions to a COA ontology for counterinsurgency operations [1, 2] to support the EC 10 demonstration.

Fig. 3 illustrates a COA outcome forecast table, which shows in stoplight format the possible OPFOR COAs. In this scenario, the CF commander uses this table to help make decisions on his own COAs based on what the OPFOR is likely to do.

Course of Action	Expected COA Outcome		
	Discourage confidence in host nation government	Impose religious and cultural control	Inflict destruction on COIN forces
Vehicle borne IED attack	green	amber	amber
Attack against CF FOB or PB	amber	amber	amber
Attack against civilian NAI	green	green	red
Sniper attack against CF PB	amber	amber	amber
Sniper attack against local official	green	amber	red
Deliberate ambush on CF patrol	green	amber	green
Hit and run against CF convoy	amber	amber	amber
Roadside IED attack in village A	red	red	green
Roadside IED attack in village B	green	green	amber

Fig. 3. The COA effect forecast table shows the expected or forecast change of state for a given COA using a red / amber / green indication

The rows in this table represent the possible OPFOR COAs, and the columns represent the assumed goals of the OPFOR. The effect of a COA for each goal is shown as red / amber / green indicators, where green indicates the best possible outcome for the OPFOR, red is the worst possible outcome the OPFOR, and amber lies somewhere in between. The indicators for each of the goals are *from the perspective of the OPFOR*. The COA effect table is created from historical data mining, manual entry by an SME or analyst, or more likely some combination of automated mining and manual entry¹.

In this particular scenario, we assume that the current state is amber for the medium- and long-term goals "discourage confidence in the host-nation government" and "establish religious and cultural control" and red for the short-term goal "inflict damage on the CF". In this state, there is room for improvement for the long-term

¹ Note that these predictions can be highly subjective and it is possible that different SMEs or analysts will come up with different forecast effects

goal and room for improvement or degradation for the medium- and short-term goals, from the perspective of the OPFOR. For example,

- If the OPFOR chooses to execute a vehicle-borne IED attack (VBIED), then it will improve its goals to inflict damage on the CF (red to amber) and discourage confidence in the host-nation government (amber to green), but will neither improve nor degrade its goal to establish religious and cultural control.
- If the OPFOR chooses to execute an attack against a CF forward operating base (FOB) or patrol base (PB), then it will improve its goal to inflict damage on the CF (red to amber) and not impact any of the other OPFOR goals.

Fig. 4 illustrates the COA preference graph, which is a representation of the preference structure of a decision maker. Each node in the graph represents a possible COA and edges represent the “is-preferred-to” relation. An edge from a source node to a target node indicates that the source node is preferred to the target node, from the perspective of a decision maker. The contents of the node show the forecast effects of the corresponding COA, as described above.

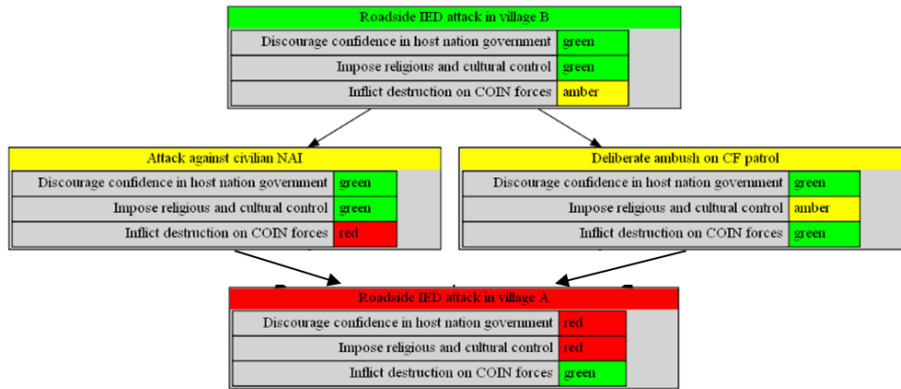


Fig. 4. The Preference Graph shows the preference structure for a collection of COAs from the perspective of a specific decision maker.

The colored title bar of each node represents the overall preference of each COA: a green title indicates that the COA is the most preferred; a red title indicates that the COA is the least preferred; and an amber title indicates that the COA is neither the most nor the least preferred. By inspection, a commander is able to visualize the most- and least-preferred COAs. Using this graph, the own force commander can assess which COA the OPFOR is likely to pursue and take that into account when formulating the blue force own COAs.

Fig. 5 shows the attributes and COA outcomes to support the EC 10 demonstration. The objectives of the OPFOR, shown in the upper left of the figure, are modeled as utility-theoretic attributes by extending the ontology described in section 2. These attributes include the short-term, medium-term and long-term goals described in section 1. The outcome of an OPFOR action, shown in the upper right of

the figure, is modeled as a decision-theory alternative. The key ontological modeling decision represented here is that the outcomes of the OPFOR actions are the outcomes over which the decisions are made. For EC 10, the decision is an assessment of which outcome is the most preferred COA for the OPFOR, given the objectives of the OPFOR decision maker.

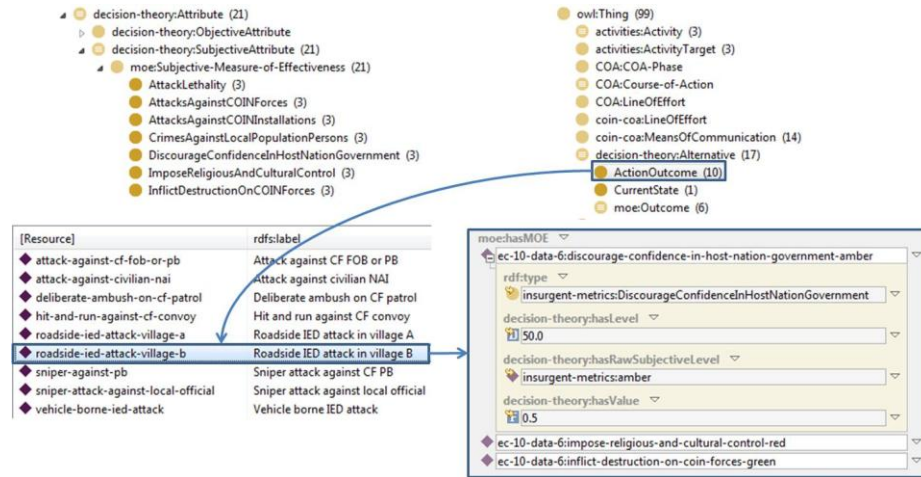


Fig. 5. Course of action ontology concepts to support the EC 10 demonstration

Individual outcomes are shown in the lower left of the figure as instances of the COA action outcome class. An example of the outcome that results from a vehicle borne IED attack (VBIED) is shown in the lower right of this figure. The outcome is described in terms of the attributes "discourage confidence in host nation government" (the medium term goal), "impose religious and cultural control" (the long-term goal) and "inflict destruction on CF forces" (the short-term goal).

The description of a COA outcome attribute is described by an attribute level, a raw subjective level, and a value. The attribute level is the source measurement of the attribute; for example, the "discourage confidence in host nation government" attribute might be measured by the rate at which the local population goes to the host-nation government to resolve legal disputes or obtain loans or other financial assistance, instead of going to the shadow insurgent government. The lower the rate, the better for the insurgents.

To support decision making in the context of the EC 10 demonstration, two rules are necessary to convert the attribute levels to a raw subjective level and to convert the raw subjective level to a utility-theoretic value.

Fig. 6 shows the SPIN² rule for assigning a raw subjective level (green / amber / red) to a given raw level band. The first and second arguments to the rule define the raw level band bounds and the third argument is the raw subjective level for those bounds. For this rule for the EC 10 demo, each of the attributes were assigned a raw

² <http://spinrdf.org>

subjective level of red for attribute levels less than 33% (in the rule, arg3 = red, arg1 = 0.0, arg2 = 33.0); green for attribute levels above 66% (in the rule, arg3 = green, arg1 = 66.0, arg2 = 100.0); and amber for all other attribute levels (in the rule, arg3 = amber, arg1 = 33.0, arg2 = 66.0). For the example shown in **Fig. 5**, the attribute level for the attribute "discourage confidence in host-nation government" for the action outcome "roadside IED attack in village A", the attribute is 50, so the rule would assign a raw subjective level of amber.

```

CONSTRUCT {
  ?this decision-theory:hasRawSubjectiveLevel ?arg3 .
}
WHERE {
  ?this decision-theory:hasLevel ?hasLevel .
  OPTIONAL {
    ?this decision-theory:hasRawSubjectiveLevel ?existingRawLevel .
  } .
  FILTER (!bound(?existingRawLevel)) .
  FILTER ((?hasLevel <= ?arg2) && (?hasLevel >= ?arg1)) .
}

```

Fig. 6. SPIN rule for assigning a raw subjective level (green, amber, red) for a given attribute level band

```

CONSTRUCT {
  ?this decision-theory:hasValue ?hasLevel .
}
WHERE {
  ?this decision-theory:hasRawSubjectiveLevel ?level .
  FILTER (?level = ?hasRawSubjectiveLevel) .
}

```

Fig. 7. SPIN rule for assigning a value ([0.0 1.0]) for a given raw subjective level

Fig. 7 shows the SPIN rule for assigning a utility value to a raw subjective level (green / amber / red). The first argument to the rule (hasLevel) is the utility value for the second argument, the raw subjective level (hasRawSubjectiveLevel). For this rule for the EC 10 demo, a raw subjective level of red was assigned a utility value 0.0 (in the rule, hasLevel = 0.0, hasRawSubjectiveLevel = red); green was assigned a utility value 1.0 (in the rule, hasLevel = 1.0, hasRawSubjectiveLevel = green); and amber was assigned a utility value 0.5 (in the rule, hasLevel = 0.5, hasRawSubjectiveLevel = amber). For the example shown in **Fig. 5**, the raw attribute level for the attribute "discourage confidence in host-nation government" for the action outcome "roadside IED attack in village A" is amber, so the rule would assign a utility value of 0.5.

4 Conclusions

Initial results from the Empire Challenge 10 demo showed promise for the approach described in this paper, especially the COA effect table and preference graph. Potential users were able to clearly assess the changes to state for each OPFOR COA as well as the most- and least-preferred COAs for the OPFOR decision maker. The COA ontology, augmented with concepts from utility theory, provides a strong theoretical foundation for creating the preference graph and using it as a COA assessment tool.

While utility theory has been used in modeling decision problems similar to COA assessment and selection, the marriage of the utility-theoretic model and semantic technologies has conferred the following benefits:

- The underlying utility theory model provides a theoretically sound foundation for defining useful properties and rules to support COA selection and assessment; these properties and rules are easily modeled using semantic technologies
- The ability to transform raw data, using a handful of simple SPARQL rules, into RDF-based representations to support visualizations that are natural for military decision makers; for example, the red / amber / green visualizations in the COA effect table.
- The ability to quickly modify the preference structure of a decision maker in a dynamic environments.

Acknowledgements: This work was supported by the Office of Naval Research under Contract N00014-09-C-0334.

5 References

1. Darr, T. P., Benjamin, P. and Mayer, R., "Course of Action Planning Ontology", Ontology for the Intelligence Community Conference (OIC 2009), George Mason University, October 21-23, 2009.
2. Darr, T. P., Benjamin, P. and Mayer, R., "Course of Action Ontology for Counterinsurgency Operations", 15th International Command and Control Research and Technology Symposium (CCRTS), Santa Monica, CA, June 22-24, 2010
3. "Stability Operations", Headquarters of the Department of the Army, Field Manual No. 3-07 (FM 3-07), October 2008.
4. "Counterinsurgency", Headquarters Department of the Army, Field Manual No. 3-24 (FM 3-07), Headquarters Marine Corps Combat Development Command Department of the Navy, Marine Corps Warfighting Publication No. 3-33.5 (MCWP 3-33.5), December 2006.
5. "Information Operations: Doctrine, Tactics, Techniques and Procedures", Headquarters Department of the Army, Field Manual No. 3-13 (FM 3-13), November 2003
6. "Empire Challenge 10", Date Accessed: August 27, 2010, Available: http://www.jfcom.mil/about/fact_ec10.html.
7. "EC 10: Green Devil provides improved tracking", Date Accessed: August 27, 2010, Available: <http://usjcom.dodlive.mil/2010/08/10/ec-10-green-devil-provides-improved-tracking/>.
8. "Empire Challenge 10 finishes up", Date Accessed: August 27, 2010, Available: <http://www.jfcom.mil/newslink/storyarchive/2010/pa081410.html>.
9. Keeney, R.L. and Raiffa, H. "Decisions with Multiple Objectives: Preferences and Value Tradeoffs", Wiley and Sons, New York, 1976